AFRL-RI-RS-TR-2019-099

ENABLING ANYCAST IN THE RESEARCH ROOT (EARR)

UNIVERSITY OF SOUTHERN CALIFORNIA / INFORMATION SCIENCES INTSTITUTE

MAY 2019

FINAL TECHNICAL REPORT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

STINFO COPY

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE

■ **AIR FORCE MATERIEL COMMAND** ■ **UNITED STATES AIR FORCE** ■ **ROME, NY 13441**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09. This report is available to the general public, including foreign nations. Copies may be obtained from the Defense Technical Information Center (DTIC) (http://www.dtic.mil).

AFRL-RI-RS-TR-2019-099 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE CHIEF ENGINEER:

/ S /
FRANCES A. ROSE
Work Unit Manager

/ S /
QING WU
Technical Advisor, Computing & Communications Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
MAY 2019	FINAL TECHNICAL REPORT	MAY 2017 – NOV 2018

4. TITLE AND SUBTITLE

ENABLING ANYCAST IN THE RESEARCH ROOT (EARR)

5a. CONTRACT NUMBER
FA8750-17-2-0096

5b. GRANT NUMBER
N/A

5c. PROGRAM ELEMENT NUMBER
N/A

6. AUTHOR(S)

John Heidemann

5d. PROJECT NUMBER
DHS0

5e. TASK NUMBER
EA

5f. WORK UNIT NUMBER
RR

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of Southern California / Information Sciences Institute
4676 Admiralty Way, Ste. 1001
Marina del Rey, CA 90292

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Air Force Research Laboratory/RITE
525 Brooks Road
Rome NY 13441-4505

10. SPONSOR/MONITOR'S ACRONYM(S)
AFRL/RI

11. SPONSOR/MONITOR'S REPORT NUMBER
AFRL-RI-RS-TR-2019-099

12. DISTRIBUTION AVAILABILITY STATEMENT
Approved for Public Release; Distribution Unlimited. This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09

13. SUPPLEMENTARY NOTES

14. ABSTRACT

This final report summarizes the objectives for the EARR project (FA8750-17-2-0096) and the technical progress made against those objectives. The research objective of the EARR project was to (1) improve our understanding of anycast and its role in defending DNS (Domain Name System) servers against Distributed-Denial-of-Service (DDoS) attacks, (2) improve anycast and instrumentation at B-Root, and (3) improve understanding of DNS leakage and privacy. EARR accomplished those objectives, producing research reports that document the role of anycast in DDoS defense, deploying anycast at B-Root for the first time (with deployment to a third site underway at end-of-contract), and supporting analysis of DNS datasets from JAS Advisors for DNS leakage and privacy, and distributing those datasets through the DHS IMPACT program.

15. SUBJECT TERMS
Anycast, Domain Name System, Internet Topology Data, Blackhole Address Space Traffic Data

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			**FRANCES A ROSE**
U	U	U	UU	24	19b. TELEPHONE NUMBER (*Include area code*) **N/A**

Table of Contents

List of Figures

1.0 SUMMARY

This report shows accomplished objectives over the course of the project. The effort improved understanding of anycast and its role in defending DNS (Domain Name System) servers against Distributed-Denial-of-Service (DDoS) attacks, as shown in several peer-reviewed technical papers and software releases such as Verfploeter. The effort also improved anycast and instrumentation at B-Root, deploying a second site at Miami (in addition to Los Angeles), with a third site underway. All sites are served by a common dashboard. Finally, the effort improved understanding of DNS leakage and privacy, working with JAS Advisors to anonymize, release, and analyze data from corp-.com

2.0 INTRODUCTION

This report provides a comprehensive, cumulative and substantive summary of the work done as part of the EARR project (FA8750-17-2-0096).

The goals of this project were to (1) improve the understanding of anycast and its role in defending DNS (Domain Name System) servers against Distributed-Denial-of-Service (DDoS) attacks, (2) improve anycast and instrumentation at B-Root, and (3) improve understanding of DNS leakage and privacy. EARR accomplished those objectives, producing research reports that document the role of anycast in DDoS defense, deploying anycast at B-Root for the first time (with deployment to a third site underway at end-of-contract), and supporting analysis of DNS datasets from JAS Advisors for DNS leakage and privacy, and distributing those datasets through the DHS IMPACT program..

This work was carried out by USC/ISI and JAS Advisors over the period of performance from May 2017 to November 2018. In this introduction we provide overall administrative information and a public description and overview of the problem and approach.

2.1 Research Objectives

The research objective of EARR was to further understand the use, operation, and evolution of the DNS. It supported the DHS S&T CSD mission to enhance “the security and resilience of the nation's critical information infrastructure and the Internet”. It was in response to area *FY16 Cyber Agility* in BAA-AFRL-RIK-2015-0015. Cyber agility is a fundamental challenge of Defensive Cyber Operations, with the need to understand what actions are necessary to ensure adversaries’ attacks are unsuccessful. This research advanced that goal in these ways:

DNS Defense and Resiliency to Attack: The DNS infrastructure today is often subject to

Distributed Denial-of-Service (DDoS) attacks. Anycast [Avramopoulos09a] plays a key role in defending against DDoS attacks, yet its deployment and operation is poorly understood, particularly under stress. Recent work has begun to examine public data to understand the behavior of anycast under stress in the DNS [Moura16a], but better methods are needed to harden deployments in the DNS to attacks, to reallocate resources during attacks, and to automate defensive measures. Data needs to be gathered from actual DNS deployments and shared with the community, and the ability to experiment on this data in a testbed is essential to develop, test, improve, and automate defensive measurements against attack.

DNS as part of the Internet Ecosystem: As a critical part of the Internet infrastructure, DNS is part of the interplay of attacks and defenses in the Internet ecosystem. The DNS is a component in attacks that exploit amplification [Kambourakis07a, Rossow14a], injection for hijacking [Herzberg13a], and interception for censorship [Duan12a, Anonymouas14a]. In addition, the DNS infrastructure forms an ecosystem of its own, with multiple providers and operators interacting in many ways. The DNS Root is one of the most diverse examples of this architecture, with 13 distinct deployments (the 13 root "letters") that are operated by 12 organizations. Traffic data, system performance metrics, and operational experience, are required to understand how these many anycast implementations perform under different conditions, and how future deployments should be designed to maximize diversity and resilience to defend against a range of threats.

Research in these methods requires the ability to collect, share, and experiment with DNS data, and to test new ideas about defenses, deployments, and new designs, all under relevant conditions. This effort enhanced the capability to provide and experiment with operational DNS data. This project made data available to the cybersecurity research community at-large via DHS S&T CSD's Information Marketplace for Policy and Analysis of Cyber-risk & Trust (IMPACT, at https://impactcybertrust.org), in cooperation with the LACREND project (https://ant.isi.edu/lacrend/).

Researchers also require the ability to experiment under near-real-world conditions. This effort developed a framework and testbed that allows third-party researchers to experiment with DNS traffic and servers in a testbed that closely mimics a real-world deployment.

Finally, this effort was an additional effort added to address additional questions. The context for these additional questions is beyond the global DNS name space: private namespace are causing namespace collisions with the global DNS, resulting in complex and pervasive occurrences that manifests throughout the global Internet namespace and have the potential to expose serious security-related issues for users of the DNS. These collisions and resulting traffic leakage potentially threaten the privacy and security of the United States' enterprises and organizations.

2.2 Public Research Goals/Contribution.

The Domain Name System (DNS) is an important part of nearly every Internet transition, so correct operation of the DNS is essential. The DNS at risk from Distributed Denial-of-Service (DDoS) attacks, and subject to stress as it grows and evolves. Better tools are needed to support analysis of this system to improve its defense and evolution. Under this effort we advanced the state of the art in tools to analyze DNS traffic.

The research objective of EARR was to improve our understanding of the use of anycast to respond to stressful events like DDoS attacks. We documented current attacks on infrastructure, improved our ability to use that data to replay attacks and evaluated multiple defensive options. In addition, we examined questions about DNS name leakage and privacy.

2.3 Expected Impact

The primary expected impact of this work was new datasets and tools supporting analysis of the Domain Name System by academic, government, and commercial researchers of

the Internet. Secondary impacts include results demonstrating ways to improve the DNS using anycast, and the release of specific datasets about the DNS.

3.0 Methods, Assumptions, and Procedures

This section summarizes the research methods, assumptions, and procedures. Detailed results follow in the next section.

3.1 Methods

The EARR project improved the understanding of the use of anycast to respond to stressful events like DDoS attacks. The effort documented attacks on DNS infrastructure, improved our ability to ability to use that data to replay attacks and evaluate defensive options. This work had one base component and three options:

1. (Base) Extending DNS measurement infrastructure to support anycast, including (a) deploying DNS measurement at a second anycast site for B-Root, and (b) developing independent and integrated views of DNS information from both sites.

2. (Option A) Develop a near-real-time dashboard for B-Root and multiple DNS root letters. (a) Develop measurement and reporting infrastructure, (b) develop a dashboard reporting status, (c) provide reports to multiple root letters.

3. (Option B) Provide research access to DNS testbed prototype. (a) Develop a framework to allow per-user access to testbed, (b) Document testbed use for a third user, (c) Pilot testbed use with 2 external researchers.

4. (Option C, EARR-RING) Evaluate interactions between the global DNS name space and private namespaces, considering potential name collisions and traffic leakage

3.1.1 Assumptions and Comparison with Current Technology.

Current state-of-the-art relevant to this project is:

1. For DNS capture: general purpose capture systems like tcpdump (http://www.tcpdump.org) capture packets but do not explicitly support DNS. DNS-specific systems such as DSC (https://www.dns-oarc.net/oarc/data/dsc) capture DNS traffic but do not provide specific services for anonymization. We propose to build on the LANDER packet capture system (https://ant.isi.edu/software/lander/) and its ability to provide multiple queues of data streams, each with different users and levels of processing (raw packets, DNS streams, etc.) and anonymization (for clear to anonymized).

2. For DNS anonymization, we plan to build on techniques such as prefix-preserving IP anonymization. We will build on the existing dag_scrubber tool built at USC/ISI; we expect to build a new DNS-specific tool.

3. For DNS replay, we have developed LDPlayer, a new DNS replay system that can interoperate with live tools (DNS servers like bind and unbound). LDPlayer components are available at https://ant.isi.edu/software/ldplayer/ .

3.1.2 Procedures: Tasks and Deliverables.

The procedures of our work are summarized by the following tasks, and deliverables. Details about the work that took place are summarized in the next section (Results and Discussion).

3.1.2.1 Detailed Individual Task Descriptions.

Progress against planned objectives is summarized below, with details about progress towards planned objectives listed in Section 4.1 marked [STx].

Approach: This effort had a base and multiple options that could be selected depending on level of support that was available. The options build on the base, and option B builds on option A.

Base: Extend measurement infrastructure to support DNS anycast.

Task 1 (base): Extend measurement architecture to support anycast.

Objective 1.1: DNS measurement at a second anycast site for B-Root. *Objective completed as of 2017q2.*

Objective 1.2: software to provide both independent (per-site) and integrated views of the DNS query stream. *Objective completed as of 2017q3.*

Option A: Develop a near-real-time dashboard and alerting system for B-Root and Multiple Root Letters. (Exercised.)

Task 2 (part of Option A): near-real-time alerting system for B-Root.

Objective 2.1: Develop base measurement system. *Objective completed as of 2018-06.*

Objective 2.2: Develop a dashboard reporting current status. *Objective completed in 2018 with [ITP4]*.

Objective 2.3: Provide alerting to other root letters.

Option B: Provide Research Access to Performance Experiments in Testbed. *(Exercised.)*

Task 3 (part of Option B): Testbed use by external researchers.

Objective 3.1: Develop a framework for per-user accounts for experiments on testbed. *Objective retargeted in 2018-08 with [ITP6] to software supporting multiple testbeds.*

Objective 3.2: Document tools: remote access, trace replay, and reporting. *Objective completed in 201810-05 with [ITP7]*.

Objective 3.3: Operate testbed on a pilot basis for at least two external researchers. *Objective completed 2018-08 with Robert Story as first non-EARR user, and in 2019 with additional users.*

Option C: Additional Work Examining Naming and Potential Collisions *(Option C was executed to begin 2017-08-01, running for 12 months.)*

Objective C.1: Research Increment on Naming -- *JAS Global Advisors LLC*

Subobjective C.1.1: Obtain technical control of a number of domains that have potential data pertaining to namespace collisions. *Objective completed 2017-12.*

Subobjective C.1.2: Create a robust technical hosting, data collection, and storage infrastructure to capture and collect data for researchers regarding the domain registrations in Result 1. *Objective completed 2018-03.*

Objective C.2: Supply Data to DNS Researchers via IMPACT -- *JAS Global Advisors LLC and USC/ISI*

Subobjective C.2.1: Release an anonymized version of the data collected from Objective 1 to qualified, vetted researchers as a part of the DHS IMPACT program. *Objective completed 2018-03.*

Objective C.3: DNS Research on Critical Domains – *USC/ISI*

Subobjective C.3.1: Study the datasets created in Objective 2 for name-leakage, timing disclosures, DDoS and other noteworthy security events along with new insights into enterprise name spaces, such as leaked domains that may be of interest to DHS and DNS researchers. *Initial study was completed in 2018-03 with ISI assistance in evaluating anonymization schemes of the JAS data. Additional evaluation is in this report as section 3.3. Objective completed 2018-11.*

3.1.3 Deliverables Description.

Deliverable 1.1: (Base) Demonstrate the new anycast measurement capability provided DNS data corresponding to the Day-in-the-Life-of-the-Internet experiments from both anycast sites, to correspond with the measurement dates selected by DNS-OARC. *Deliverable completed as of 2017-08.*

Deliverable 1.2: (Base) Provided the software developed as part of Objective 1.2 as open source software, to be distributed on a public website. *Deliverable completed 2017-09-14 as noted in [ITP5], software provided as https://ant.isi.edu/software/verfploeter/.*

Deliverable 2.1: (Option A) Provided documentation for the alerting system. *Deliverable completed with this final report.*

Deliverable 3.1: (Option B) Provides documentation for testbed use suitable for supporting the external users. *Deliverable completed 2018-10-05 as noted in [ITP6]. Documentation is provided in the technical paper [Pub5] and at the website https://ant.isi.edu/software/verfploeter/*

Deliverable 4.1: (Option C) Provisioned EARR-RING-related datasets in IMPACT. *Deliverable completed 2018-03 as noted in [IPT3]. For data, see https://impactcybertrust.org/.*

4 RESULTS AND DISCUSSION

This section provides a detailed discussion of results of the project. We begin with a narrative of accomplishments, including an enumeration of improvements to prototypes and technology-transfer actions. The report then highlights publications and presentations (both opportunities for technology transfer). Finally, the report summarizes key results and provides brief summarizes of additional results.

4.1 Technical Accomplishments Over Contract

Accomplishments below are identified as pertaining to task objective [OT1.x], deliverable [D1.x], Improvements To Prototypes [ITP], technology transfer [TTx], presentation [PreX], and publication [PubX] (where the *x* in each of these indicates which sequence that accomplishment is in that category).

This report includes discussion about some activities related to B-Root that are not directly funded by EARR. These activities are identified as [ext] and are provided to give context for the EARR work. They are supported by other (non-EARR) funding.

1. **[OT1.1][ITP1]** As of 2017-08, B-Root has deployed DNS service and data collection at its Miami location, completing anycast. *Completes Deliverable 1.1*

2. **[TT1] [Pre1] [ext]**On 2017-06-19 John Heidemann gave the invited keynote "Digging in to Ground Truth in Network Measurements" at the IEEE TMA PhD School in Dublin, Ireland. This talk included discussion of DDoS events related to EARR. This work was not directly supported by EARR.

3. **[TT2] [Pub2] [ext]** John Heidemann, with Kensuke Fukuda and Abdul Qadeer, published "Detecting Malicious Activity with DNS Backscatter Over Time". This work pre-dates EARR and was not directly supported by EARR, but it uses data collected as part of EARR.

4. **[TT3] [Pub3][Pub4] [ext]** Liang Zhu, working with John Heidemann, has been developing LDplayer, a DNS trace replace system that we expect to support Task 3. This work resulted in a poster that we released as a technical report. This task is not currently selected as part of EARR and Zhu's work is currently supported by other projects.

5. **[TT4] [Pub1][Pre2] [ext]** Lan Wei, working with John Heidemann, has been evaluating anycast. This work resulted in a publication and presentation in 2017q3. Although this work is not directly supported by EARR, it contributes to understanding of anycast in the B-Root infrastructure.

6. [TT5][Pub5]Wouter B. de Vries, Ricardo de O. Schmidt, Wes Hardaker, John Heidemann, PieterTjerk de Boer and Aiko Pras 2017. **Verfploeter: Broad and Load-Aware Anycast Mapping**. *Proceedings of the ACM Internet Measurement Conference* (London, UK, 2017), 477– 488. <https://www.isi.edu/%7ejohnh/PAPERS/Vries17b.html>. It was not directly supported by EARR, but it contributes to understanding of anycast in the B-Root infrastructure.

7. [TT6][Pre4] On 2017-09-30, Wes Hardaker gave a talk "Verfploeter: Broad and Load-Aware Anycast Mapping" at DNS-OARC (a second time). This paper is joint work of Heidemann and researchers at U. Twente and SIDN Labs. It was not directly supported by EARR, but it contributes to the B-Root infrastructure.

8. [ITP2] In 2017-09, Wes Hardaker deployed a prototype of his localroot at https://localroot.isi.edu.

9. [TT7][Pre5] On 2017-11-11, Wes Hardaker gave a talk "Verfploeter: Broad and Load-Aware Anycast Mapping" at MAPRG (a third time) with an emphasis on how the Verfploeter technique can help routing protocol design. It was not directly supported by EARR, but it contributes to the B-Root infrastructure.

10. [TT7] [Pre6] On 2017-11-15, Wes Hardaker gave the talk "LocalRoot – Serve Yourself" at the DNSSEC Workshop.

11. As of 2017-11, The EARR-RING contract extension is underway but not yet fully executed. We have added its deliverables in this report, but will only begin work in the next reporting period assuming the contract is executed.

12. In 2017-12 we completed contracting for EARR-RING and welcomed JAS Advisors as a subcontractor. This report therefore adds Option C with corresponding tasks and deliverables.

13. [TT8][Pre7] On 2017-12-07, John Heidemann and Christos Papadopoulos gave the talk “Los Angeles/Colorado Application and Network Information Community (LACANIC): Project Update, Dec. 2017” at the IMPACT PI Meeting. This presentation included some information about EARR

14. [TT9] [Pre8]On 2017-12-07, Jeff Schmidt presented “Introducing: The ORDINAL Dataset (Operational Research Data from Internet Namespace Logs” at the IMPACT PI meeting.

15. [OT.C.1.1] In 2017-12, JAS Advisors completed technical control of the domains under consideration as part of objective C.1.1.

16. [OT.C.1.2][OT.C.2.1][OT.C.3.1] In 2018-02, JAS Advisors and ISI began discussions about the different datasets in ORDINAL, what can be released, and what kinds of anonymization is required.

17. [TT10][Pre9] On 2018-02-18, Wes Hardaker gave a talk “Analyzing and Mitigating Privacy with the DNS Root Service”. It was not directly supported by EARR but made use of B-Root infrastructure enhanced by EARR efforts.

18. [TT11][Pre10] On 2018-02-18, Basileal Imana gave the talk “Enumerating Privacy Leaks in DNS Data Collected above the Recursive” at the IEEE NDSS DNS Privacy Workshop in San Diego, California. This talk was based on the paper of the same name by him and Aleksandra Korolova and John Heidemann.

19. [TT12][Pre11] On 2018-02-18, Wes Hardaker gave the talk “Analyzing and Mitigating Privacy with the DNS Root Service” at the IEEE NDSS DNS Privacy Workshop in San Diego, California. This talk was based on his paper of the same name.

20. [OT.C.2.1] [ITP3] In 2018-03, JAS Advisors completed steps to prepare anonymized datasets and provisioned those datasets into the IMPACT portal, completing an initial version of task deliverable [OT.C.2.1]. This dataset used a modified version of the dnsanon tool provided by USC/ISI.

21. [TT14][Pre13]On 2018-03-09. Wes Hardaker gave the talk “Analyzing DITL Data From BRoot” at the DNS-OARC meeting.

22. [TT13][Pre14] On 2018-03-11, Wes Hardaker gave the talk “LocalRoot—Serve Yourself” at the ICANN Technology Day in at San Juan, Puerto Rico.

23. [TT15][Pre15] On 2018-04-05, Basileal Imana gave the talk “Enumerating Privacy Leaks in DNS Data Collected above the Recursive” at the ISI Graduate Student Symposium in Marina del Rey, California. This talk was based on the IEEE NDSS DNS Privacy Workshop paper of the same name by him and Aleksandra Korolova and John Heidemann

24. [TT17][Pre16] On 2018-05-15, Jeff Schmidt gave the talk "Everything You Know About the Domain Name System (DNS) Is Wrong" at Ohio ISSA. This talk referenced and marketed IMPACT datasets.

25. [TT18][Pre17] On 2018-06-15, Jeff Schmidt gave the talk "The Saga and Intrigue of CVE-20150008/corp.com" at 614con. This talk referenced and marketed IMPACT datasets [TT10][Pub7] On 2018-05-30, we released the following technical report: Giovane C. M. Moura, John Heidemann, Moritz Müller, Ricardo de O. Schmidt, and Marco Davids. When the Dike Breaks: Dissecting DNS Defenses During DDoS (extended). Technical Report N. ISI-TR-725, USC/Information Sciences Institute, May, 2018. <https://www.isi.edu/%7ejohnh/PAPERS/Moura18a.html>.

26. [ITP4][TT19][OT1.1]. In 2018 we have deployed an integrated dashboard that includes data from multiple sites in our Grafana console. This deployment completes Task 1.2.

27. [TT20][OT2.1] In 2018 we have deployed a preliminary alerting system for B-Root based on Nagios. This system is currently in use internally in B-Root. This deployment completes task 2.1

28. [Pub8][Pub9][Pub10]In 2018-08 we had three papers accepted to appear in ACM Internet Measurement Conference 2018. We prepared camera-ready papers to present in the next reporting period.

29. [Pre18] On 2018-06-25, Wes Hardaker gave the talk “RFC8145 / KSK-2010 Signal Analysis” at the DNSSEC Workshop at ICANN62 in Panama City, Panama. This talk built on analysis of EARR-derived data, and it influenced the ICANN policy on the DNSSEC key roll planned for 2017.

30. [Pre19] On 2018-07-15, John Heidemann gave the talk “When the Dike Breaks: Dissecting DNS Defenses for DDoS” to Root-Ops in Montreal, Canada. (John Heidemann gave the talk remotely from Los Angeles.) This talk was on a paper that will be presented at IMC 2018.

31. [Pre20] On 2018-07-16, Wes Hardaker gave the talk “draft-dns-zone-digest” at the DNSOP working group at IETF 102 in Montreal, Canada. Although not directly supported by EARR, this work discussed DNS related to EARR.

32. [Pre21] On 2018-07-19, Wes Hardaker gave the talk “RFC8145 / KSK-2010 Signal Analysis” at the IRTF MAPRG research group at IETF 102 in Montreal, Canada This talk built on analysis of EARR-derived data, and it influenced the ICANN policy on the DNSSEC key roll planned for 2018.

33. [D1.2][ITP5] Although not previously reported, on 2017-09-14, USC released the Verfploeter software used to evaluate anycast catchments at https://ant.isi.edu/software/verfploter/ *This software release completes deliverable D1.2.*

34. [O3.1] In 2018-08, we decided that rather than target a single testbed that supports remote user access, we are focusing on releasing trace replay software that allows *anyone* to operate their own testbed, plus datasets to allow recreation of specific scenarios.

35. [TT][O.3.3] In 2018-08, Robert Story used a pre-release version of LDPlayer to carry out performance evaluation of B-root. He represents one external (non-EARR) user of our DNS testbed and trace replay software. Robert compared LDplayer to traditional DNS load-testing tools, finding that binary input is required to get good performance out of LDPlayer, and evaluating the performance under stress of a new hardware design for B-

Root. We plan to use this new hardware design in the third B-Root site, with deployment expected in 2018q4.

36. [TT] As of 2018-10, Wouter de Vries (graduate student at U. Twente and co-developer of Verfploeter) has completed a summer internship at Cloudflare. While there, he worked with them to deploy Verfploeter inside a commercial DNS's anycast infrastructure, and this tool has since gone live. This accomplishment represents a significant example of technology transfer of technology supported by EARR.

37. [D3.1][ITP7][O3.2] On 2018-10-05, Liang Zhu released LDPlayer, a platform for DNS trace replay, mutation, and experimentation, at https://ant.isi.edu/sofwtware/ldplayer/ this software release completes deliverable D3.1.

38. [TT][O.3.3] In 2018-10, A.S.M. Rizvi carried out experiments using LDplayer where he evaluated several DDoS and traffic burst scenarios against his automated DDoS defenses. He represents a second external (non-EARR) user of our testbed and trace replay software and datasets. He has curated datasets representing the events of interest and we plan to release them to the public in 2018q4.

39. [TT][O.C.3.1] In 2018-11, Wes Hardaker did additional evaluation of the JAS corp.com data. Results of that evaluation are reported in section 4.6 of this report. This evaluation, with prior work with JAS evaluating their data for anonymization, completes Objective C.3.1.

4.2 Significant Changes to Technical Approach Over Project

One significant change to technical approach from what was proposed is to add Options A, B, and C at different times.

In 2018-08, indicated by [IPT6] in the timeline, in evaluating the testbed, we decided to focus on software that allows trace replay combined with datasets that allow recreation of specific scenarios such as DDoS events. We focused on this combination instead of a single testbed with remote access, since it will allow the creation of *many* testbeds that can be used concurrently.

4.3 Technology Transition Outcomes

The primary technology transition plan was (1) to work with the B-Root operators to deploy developed technology in production use. (2) to work with the LACREND project to provide tools and developed datasets for distribution to the IMPACT program.

Where possible the effort released tools as open source software via the ANT website (https://ant.isi.edu/).

The effort documented release of Verfploeter, as described in deliverables, above. The effort has also documented the release of several datasets through IMPACT:

Dataset	Compressed
B-Root_Load-20170515	46M
B-Root_Load-20170412	46M
B_Root_Anomaly-20170425	315G
B_Root_Anomaly-20170306	727G
B_Root_Anomaly-20170221	361G
B_Root_Anomaly-20151130	552G
B_Root_Anomaly_message_question-20170425	129G
B_Root_Anomaly_message_question-20170306	310G
B_Root_Anomaly_message_question-20170221	155G
B_Root_Anomaly_message_question-20151130	287G
B_Root_week_message_question-20190109	tbd
DITL_B_Root-20180410	535G
DITL_B_Root-20170919	794G
DITL_B_Root-20170411	487G
DITL_B_Root-20161001	575G
DITL_B_Root-20160405	412G
DITL_B_Root-20150413	490G
DITL_B_Root-20140428	169G
DITL_B_Root-20130528	281G
DITL_B_Root_message_question-20180410	225G
DITL_B_Root_message_question-20170919	337G
DITL_B_Root_message_question-20170411	213G
DITL_B_Root_message_question-20161001	309G
DITL_B_Root_message_question-20160405	210G
DITL_B_Root_message_question-20150413	150G
DITL_B_Root_message_question-20140428	52G
DITL_B_Root_message_question-20130528	89G
Tangled_Verfploeter-20170323	tbd
Tangled_Verfploeter-20170201	tbd
B_Root_Anomaly-20170221	310.7G
B_Root_Anomaly-20170306	726.0G
B_Root_Anomaly-20170425	315.8G

4.4 Key Results: Detailed Summary of Analysis of JAS Corp.Com Data

USC/ISI has applied some of the DNS analysis and anomaly detection tools to the *corp.com* DNS log data. Specifically, we searched for trend shifts in DNS names and source addresses sending data to the *corp.com* name servers. The tools quickly detected a number events relating to upticks in Kerberos logins, LDAP usage, leaking Microsoft Windows domain controller requests, outcoming and incoming SMTP (E-Mail) events, and general leakage of companies names with misconfigured systems. We've shown the datasets to contain valuable research data and are interested in exploring them in greater depth in the future, looking for weak signals buried within them besides the more obvious to detect "peaks".

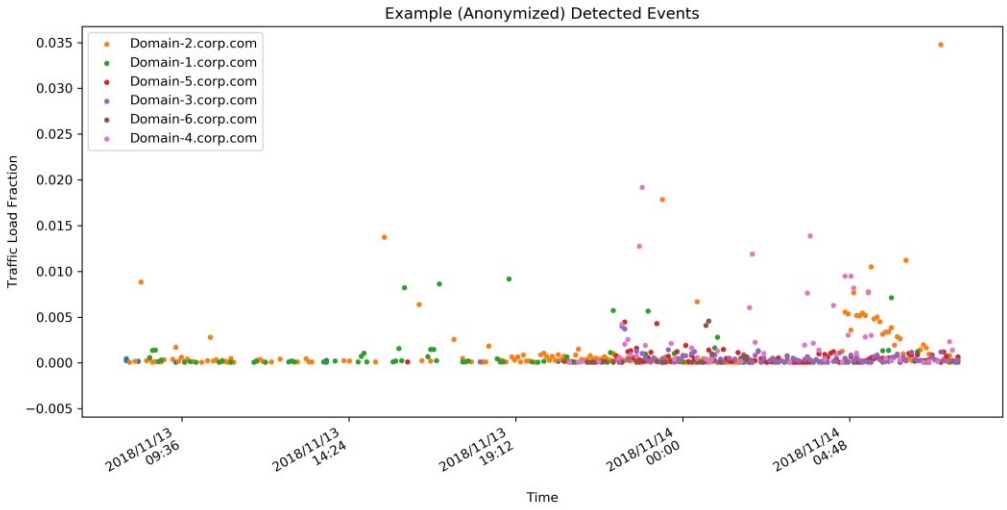

Figure 1. Detected Events

This analysis, Figure 1., establishes that are some interesting events embedded in the JAS-provided corp.com data.

4.5 Key Results: Detailed Summary of B-Root Dashboard

A screenshot of the B-Root dashboard, Figure 2, can be seen below. (Please note we have redacted some operational details.)

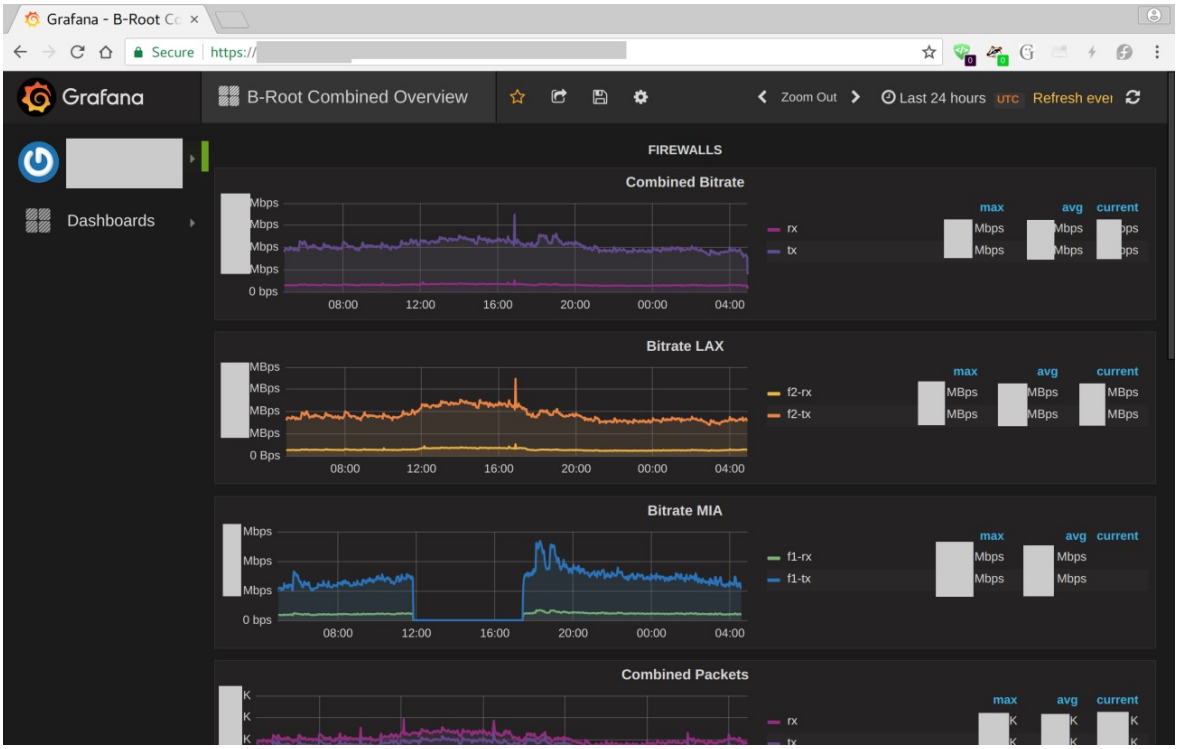

Figure 2 B-root Dashboard.

This dashboard shows the combined bitrate (top) and the bitrates for each site (middle and bottom), with both inbound and outbound traffic. This dashboard shows the two sites (Los Angeles LAX and Miami MIA) that were active when the screenshot was taken. In addition, we see that the MIA site was shut down for about seven hours for planned maintenance.

This kind of dashboard has been very useful for B-Root operation and for understanding the role of anycast in B-Root and in general.

This dashboard was also useful at providing a B-Root-specific view of the ICANN 2018 KSK Rollover, described below. Screenshots from it appeared in Wes Hardaker's talk "The KSK roll: From Concern to Calm(ish)" give at the ICANN meeting on 2018-10-24.

4.6 Key Results: Brief Summary of Multi-Letter Dashboard

This effort also worked with Verisign, ICANN, and other root operators to provide a multi-letter dashboard representing traffic across several Root Service Operators ("letters"). The effort could not reproduce that dashboard in this public report because the data is considered private by some operators.

This dashboard was of particular utility during the 2018 change of the root signing key (the KSK rollover). ICANN published a post-action report on the KSK rollover in "Review of the 2018 DNSSEC KSK Rollover" at https://www.icann.org/en/system/files/files/review-2018-dnssec-ksk-rollover-04mar19en.pdf

In addition to working with other RSOs to share data for this event, ISI analysis of B-Root data was helpful at identifying particular problems before the KSK Rollover. In particular, Wes Hardaker identified a VPN software vendor as a major source of use of the old key. He worked with them to identify this problem, and they corrected it in software updates to their product. Those results can be seen in the following graph, Figure 3:

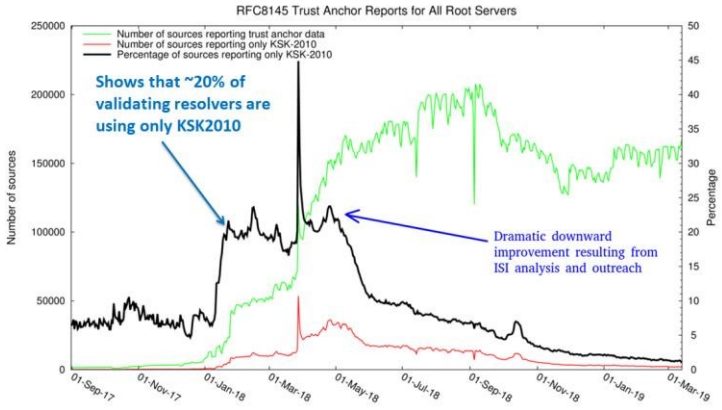

Figure 3: Impact of this effort

This work was presented at the talk “RFC8145 / KSK-2010 Signal Analysis” by Wes Hardaker at the ICANN Global Summit in June 2018 and at the MAPRG Meeting at the IETF 102 meeting in July 2018.

5 CONCLUSIONS AND RECOMMENDATIONS

5.1 Conclusions

This effort has shown that the effort accomplished the objectives over the course of the EARR project (FA875017-2-0096). It (1) improved our understanding of anycast and its role in defending DNS (Domain Name System) servers against Distributed-Denial-of-Service (DDoS) attacks, as shown in several peer-reviewed technical papers and software releases such as Verfploeter. It also (2) improved anycast and instrumentation at B-Root, deploying a second site at Miami (in addition to Los Angeles), with a third site underway. All sites are served by a common dashboard. Finally, it (3) improved understanding of DNS leakage and privacy, working with JAS Advisors to anonymize, release, and analyze data from corp.com.

5.2 Recommendations

This work has shown the importance of anycast in robust DNS operation, and has increased public awareness of this importance through technical papers. It has also improved B-Root operations, helping the Internet community.

We recommend continued and increased use of anycast, and are currently working to deploy a third anycast site for B-Root.

We also recommend continued evaluation of the role anycast serves in defending against Distributed Denial-of-Service attack. Although we have identified that role, work is ongoing to state clearly what specific actions an operator should take under different circumstances.

We recommend additional research analyzing name collisions resulting from misconfigurations and corporate networks. Our work in EARR-RING only touched the tip of the iceberg of work that we have shown exists in this area. The ICANN community is currently studying the name collision space for .corp, home and .mail but is otherwise not going to be a comprehensive study either.

6.0 REFERENCES

Lan Wei and John Heidemann. Does Anycast Hang up on You?. In *IEEE International Workshop on Traffic Monitoring and Analysis*, p. to appear. Dublin, Ireland, IEEE. July, 2017. <http://www.isi.edu/%7ejohnh/PAPERS/Wei17b.html>. It was not directly supported by EARR, but it contributes to understanding of anycast in the B-Root infrastructure.

Anycast-based services today are widely used commercially, with several major providers serving thousands of important websites. However, to our knowledge, there has been only limited study of how often anycast fails because routing changes interrupt connections between users and their current anycast site. While the commercial success of anycast CDNs means anycast usually work well, do some users end up shut out of anycast? In this paper we examine data from more than 9000 geographically distributed vantage points (VPs) to 11 anycast services to evaluate this question. Our contribution is the analysis of this data to provide the first quantification of this problem, and to explore where and why it occurs. We see that about 1% of VPs are *anycast unstable*, reaching a different anycast site frequently (sometimes every query). Flips back and forth between two sites in 10 seconds are observed in selected experiments for given service and VPs. Moreover, we show that anycast instability is *persistent* for some VPs—a few VPs never see a stable connections to certain anycast services during a week or even longer. The vast majority of VPs only saw unstable routing towards one or two services instead of instability with all services, suggesting the cause of the instability lies somewhere in the path to the anycast sites. Finally, we point out that for highlyunstable VPs, their probability to hit a given site is constant, which means the flipping are happening at a fine granularity—per packet level, suggesting load balancing might be the cause to anycast routing flipping. Our findings confirm the common wisdom that anycast almost always works well, but provide evidence that a small number of locations in the Internet where specific anycast services are never stable.

Kensuke Fukuda, John Heidemann, and Abdul Qadeer. Detecting Malicious Activity with DNS Backscatter Over Time. *ACM/IEEE Transactions on Networking*, V. 25 (N. 5), pp. 3203-3218, August 2017. <http://dx.doi.org/10.1109/TNET.2017.2724506>, <https://www.isi.edu/%7ejohnh/PAPERS/Fukuda17a.html>. This work was not directly supported by EARR, but it uses data collected as part of EARR.

Network-wide activity is when one computer (the *originator*) touches many others (the *targets*). Motives for activity may be benign (mailing lists, CDNs, and research scanning), malicious (spammers and scanners for security vulnerabilities), or perhaps indeterminate (ad trackers). Knowledge of malicious activity may help anticipate attacks, and understanding benign activity may set a baseline or characterize growth. This paper identifies *DNS backscatter* as a new source of information about network-wide activity. Backscatter is the reverse DNS queries caused when targets or middleboxes automatically look up the domain name of the originator. Queries are visible to the authoritative DNS servers that handle reverse DNS. While the fraction of backscatter they see depends on the server's location in the DNS hierarchy, we show that activity that touches many targets

appear even in sampled observations. We use information about the queries to classify originator activity using machine-learning. Our algorithm has reasonable accuracy and precision (70– 80%) as shown by data from three different organizations operating DNS servers at the root or country-level. Using this technique we examine nine months of activity from one authority to identify trends in scanning, identifying bursts corresponding to Heartbleed and broad and continuous scanning of sash.

Wouter B. de Vries, Ricardo de O. Schmidt, Wes Hardaker, John Heidemann, Pieter-Tjerk de Boer and Aiko Pras 2017. **Verfploeter: Broad and Load-Aware Anycast Mapping**. *Proceedings of the ACM Internet Measurement Conference* (London, UK, 2017), 477–488. <https://www.isi.edu/%7ejohnh/PAPERS/Vries17b.html>. It was not directly supported by EARR, but it contributes to understanding of anycast in the B-Root infrastructure.

IP anycast provides DNS operators and CDNs with automatic fail-over and reduced latency by breaking the Internet into *catchments*, each served by a different anycast site. Unfortunately, *understanding* and *predicting* changes to catchments as anycast sites are added or removed has been challenging. Current tools such as RIPE Atlas or commercial equivalents map from thousands of vantage points (VPs), but their coverage can be inconsistent around the globe. This paper proposes *Verfploeter*, a new method that maps anycast catchments using active probing. Verfploeter provides around 3.8M passive VPs, 430x the 9k physical VPs in RIPE Atlas, providing coverage of the vast majority of networks around the globe. We then add load information from prior service logs to provide calibrated predictions of anycast changes. Verfploeter has been used to evaluate the new anycast deployment for B-Root, and we also report its use of a nine-site anycast testbed. We show that the greater coverage made possible by Verfploeter's active probing is necessary to see routing differences in regions that have sparse coverage from RIPE Atlas, like South America and China.

Moritz Müller, Giovane C. M. Moura, Ricardo de O. Schmidt and John Heidemann 2017. **Recursives in the Wild: Engineering Authoritative DNS Servers**. *Proceedings of the ACM Internet Measurement Conference* (London, UK, 2017), 489-495. <https://www.isi.edu/%7ejohnh/PAPERS/Mueller17b.html>. It was not directly supported by EARR, but it uses data collected as part of EARR.

In Internet Domain Name System (DNS), services operate *authoritative* name servers that individuals query through *recursive* resolvers. Operators strive to provide reliability by operating multiple name servers (NS), each on a separate IP address, and by using IP anycast to allow NSes to provide service from many physical locations. To meet their goals of minimizing latency and balancing load across NSes and anycast, operators need to know how recursive resolvers select an NS, and how that interacts with their NS deployments. Prior work has shown some recursives search for low latency, while others pick an NS at random or round robin, but did not examine how prevalent each choice was. This paper provides the first analysis of how recursives select between name servers in the wild, and from that we provide guidance to operators how to engineer their name servers

to reach their goals. We conclude that all NSes need to be equally strong and therefore we recommend deploying IP anycast at every single authoritative.

Kensuke Fukuda and John Heidemann. Who Knocks at the IPv6 Door? Detecting IPv6 Scanning. *Proceedings of the ACM Internet Measurement Conference* (2018, Oct. 2018). It was not directly supported by EARR, but it uses data collected as part of EARR. Will be published in next reporting period.

DNS backscatter detects internet-wide activity by looking for common reverse DNS lookups at authoritative DNS servers that are high in the DNS hierarchy. Both DNS backscatter and monitoring unused address space (darknets or network telescopes) can detect scanning in IPv4, but with IPv6's vastly larger address space, darknets become much less effective. This paper shows how to adapt DNS backscatter to IPv6. IPv6 requires new classification rules, but these reveal large network services, from cloud providers and CDNs to specific services such as NTP and mail. DNS backscatter also identifies router interfaces suggesting traceroute-based topology studies. We identify 16 scanners per week from DNS backscatter using observations from the B-root DNS server, with confirmation from backbone traffic observations or blacklists. After eliminating benign services, we classify another 95 originators in DNS backscatter as potential abuse. Our work also confirms that IPv6 appears to be less carefully monitored than IPv4.

Liang Zhu and John Heidemann. LDplayer: DNS Experimentation at Scale. *Proceedings of the ACM Internet Measurement Conference* (Boston, Massachusetts, USA, Oct. 2018). It was not directly supported by EARR, but it uses data collected as part of EARR.

DNS has evolved over the last 20 years, improving in security and privacy and broadening the kinds of applications it supports. However, this evolution has been slowed by the large installed base and the wide range of implementations. The impact of changes is difficult to model due to complex interactions between DNS optimizations, caching, and distributed operation. We suggest that *experimentation* at scale is needed to evaluate changes and facilitate DNS evolution. This paper presents LDplayer, a configurable, general-purpose DNS experimental framework that enables DNS experiments to scale in several dimensions: many zones, multiple levels of DNS hierarchy, high query rates, and diverse query sources. LDplayer provides high fidelity experiments while meeting these requirements through its distributed DNS query replay system, methods to rebuild the relevant DNS hierarchy from traces, and efficient emulation of this hierarchy on minimal hardware. We show that a single DNS server can correctly emulate multiple independent levels of the DNS hierarchy while providing correct responses as if they were independent. We validate that our system can replay a DNS root traffic with tiny error (+/- 8 ms quartiles in query timing and +/-0.1% difference in query rate). We show that our system can replay queries at 87k queries/s while using only one CPU, more than twice of a normal DNS Root traffic rate. LDplayer's trace replay has the unique ability to evaluate important design questions with confidence that we capture the interplay of caching, timeouts, and resource constraints. As an example, we demonstrate the memory requirements of a DNS root server with all traffic running over TCP and TLS, and identify performance discontinuities in latency as a function of client RTT.

Giovane C. M. Moura, John Heidemann, Moritz Müller, Ricardo de O. Schmidt and Marco Davids. When the Dike Breaks: Dissecting DNS Defenses During DDoS. *Proceedings of the ACM Internet Measurement Conference* (Oct. 2018). It was not directly supported by EARR, but it uses data collected as part of EARR.

The Internet's Domain Name System (DNS) is a frequent target of Distributed Denial-of-Service (DDoS) attacks, but such attacks have had very different outcomes—some attacks have disabled major public websites, while the external effects of other attacks have been minimal. While on one hand the DNS protocol is relatively simple, the *system* has many moving parts, with multiple levels of caching and retries and replicated servers. This paper uses controlled experiments to examine how these mechanisms affect DNS resilience and latency, exploring both the client side's DNS *user experience*, and server-side traffic. We find that, for about 30% of clients, caching is not effective. However, when caches are full they allow about half of clients to ride out server outages that last less than cache lifetimes, Caching and retries together allow up to half of the clients to tolerate DDoS attacks longer than cache lifetimes, with 90% query loss, and almost all clients to tolerate attacks resulting in 50% packet loss. While clients may get service during an attack, tail-latency increases for clients. For servers, retries during DDoS attacks increase normal traffic up to 8x. Our findings about caching and retries help explain why users see service outages from some real-world DDoS events, but minimal visible effects from others.

LIST OF ACRONYMS

ACM: Association for Computing Machinery

AFRL: Air Force Research Laboratory

ANSI: American National Standards Institute

ANT: Analysis of Network Traffic Laboratory, https://ant.isi.edu/

CDN: Content Delivery Network

CPU: Central Processing Unit

CSD: Computer Science Department

CVE: Common Vulnerability and Exposure

DANE: DNS-Based Authentication of Named Entities

SMTP: Simple Mail Transport Protocol

DDoS: Distributed Denial-of-Service (attack)

DHS: Department of Homeland Security

DITL: Day-In-The-Life (of the Internet), a DNS data collection effort

DNS: Domain Name System

DNS-OARC: The DNS Operations, Analysis, and Research Center, https://www.dns-oarc.net/

DNSOP: The DNS Operations Working Group in the IETF

DNSSEC: the DNS Security protocol

DSC: DNS Statistics Collector (software), https://github.com/DNS-OARC/dsc

EARR: Enabling Anycast in the Research Root (this project)

EARR-RING: Enabling Anycast in the Research Root, Research Increment on NaminG, an extension to this project

ICANN: Internet Corporation for Assigned Names and Numbers

IEEE: Institute of Electrical and Electronics Engineers, a professional society

IETF: Internet Engineering Task Force

IMC: the ACM Internet Measurements Conference

IMPACT: Information Marketplace For Policy and ANalysis of Cyber-risk and Trust, a DHS program for data sharing, https://www.impactcybertrust.org/

IPT: Improvements to ProtoType

IRTF: Internet Research Task Force

ISI: Information Sciences Institute, a research laboratory at USC

ISSA: Information Security Summit Association

JAS: Jeff Schmidt Global Advisors, a private company

KSK: Key-Signing Key, part of DNSSEC

LACANIC: Los Angeles/Colorado Application and Network Information Community

LACREND: Los Angeles/Colorado Research Exchange for Network Data

LANDER: Los Angeles Network Data Exchange and Repository

LAX: the international airport code for Los Angeles, also the designation for a B-Root site in the Los Angeles-area

LDPlayer: Liang's DNS Player, a software package for DNS replay

LLC: Limited Liability Corporation

MAPRG: the IRTF Measurement and Analysis for Protocols Research Group

MIA: the airport code for Miami, also the designation for a B-Root site in the Miami-area

NDSS: the Network and Distributed System Security Symposium, https://www.ndss-symposium.org/

NTP: the Network Time Protocol

OMB: Office of Management and Budget

OT: Objective Task

RFC: IETF Request For Comments

RIPE: Regional Internet Registry for Europe, https://www.ripe.net

RSO: (DNS) Root Service Operator

RTT: Round Trip Time

SIDN: *Stichting Internet Domeinnaamregistratie Nederland* (Dutch: Netherlands Foundation for Internet Domain Names)

SIGCOMM: the ACM Special Interest Group for Communications

SMTP: Simple Mail Transport Protocol

TCP: Transmission Control Protocol

TLS: Transport Layer Security

TMA: The Conference on Traffic Measurement and Analysis

TT: Technical Transfer activity

USA: United States of America

USC: University of Southern California

USC/ISI: University of Southern California/Information Sciences Institute

VPN: Virtual Private Network